

Fabrizio Luccio

Mathematical issues in network construction and security

Dottorato 08

2. The graph structure of Internet, WWW, and DNS

Internet: birth and development

Born in the Boston area around 1969 as APRANET, with a starting group of four participants: UCLA, UCSB, SRI, and the University of Utah.

Since the very beginning the peculiarity of the net was the adoption of packet switching, as opposed to circuit switching used in the telephone systems.

At the beginning the net was totally financed with USA public money; in particular ARPA and NSF provided most of the funds. Fortunately the initial military purposes of use were paired with an increasing interest of the scientific community. Many people participated freely to the development of Internet with an open source mentality.

Internet: birth and development

All purposes of “purity” disappeared at once in 1995 when NSF decided to stop financing the Internet backbone, that is its central fast and high-capacity communication channels in the USA. The traffic was then re-routed on giant commercial networks that had been developed independently.

Now all the major communication companies worldwide are part of the Internet structure, and a huge number of subnets of all possible sizes participate into it. There is a substantial freedom in linking to the net and contributing to its growth.

China, now the world's biggest community of Internet users, is the main exception to adopting the Internet philosophy. Its huge new three-level national network, built with very advanced technology, is totally controlled by the government.

Internet: basic concepts and terminology

Tanenbaum. "Internet is not at all a network, but a vast collection of different networks that use common protocols and provide common services".

The vertices (nodes) of the Internet graph are computers with the functions of hosts, servers, and routers. These functions are not necessarily disjoint.

Hosts are (broadly speaking) the computers of users.

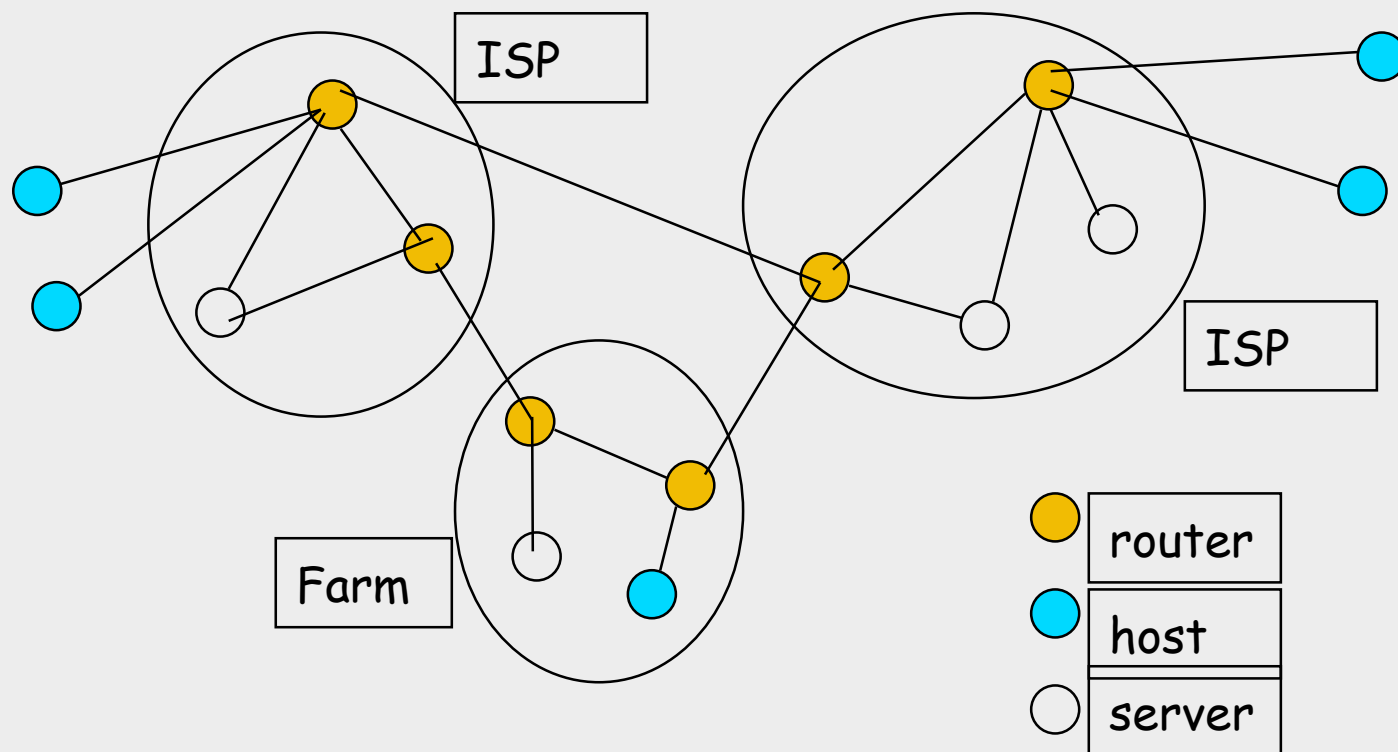
Servers provide network services.

Routers control network traffic.

The exchange of data between nodes is governed by a protocol called TCP/IP (Transmission Control Protocol/Internet Protocol), that is peculiar of Internet. It has the role of transporting the data between nodes (TCP), and controlling the routing, i.e. the next node to reach at each step (IP).

Internet as a graph

In addition to single hosts in the periphery, the nodes are arranged in independent sub-networks called Autonomous Systems (AS), interconnected through their routers. AS's may be Internet Service Providers (ISP) whose routers direct the network traffic for a set of clients, or Farms which are in general big institutions or companies that have their own net of servers, routers, and hosts.



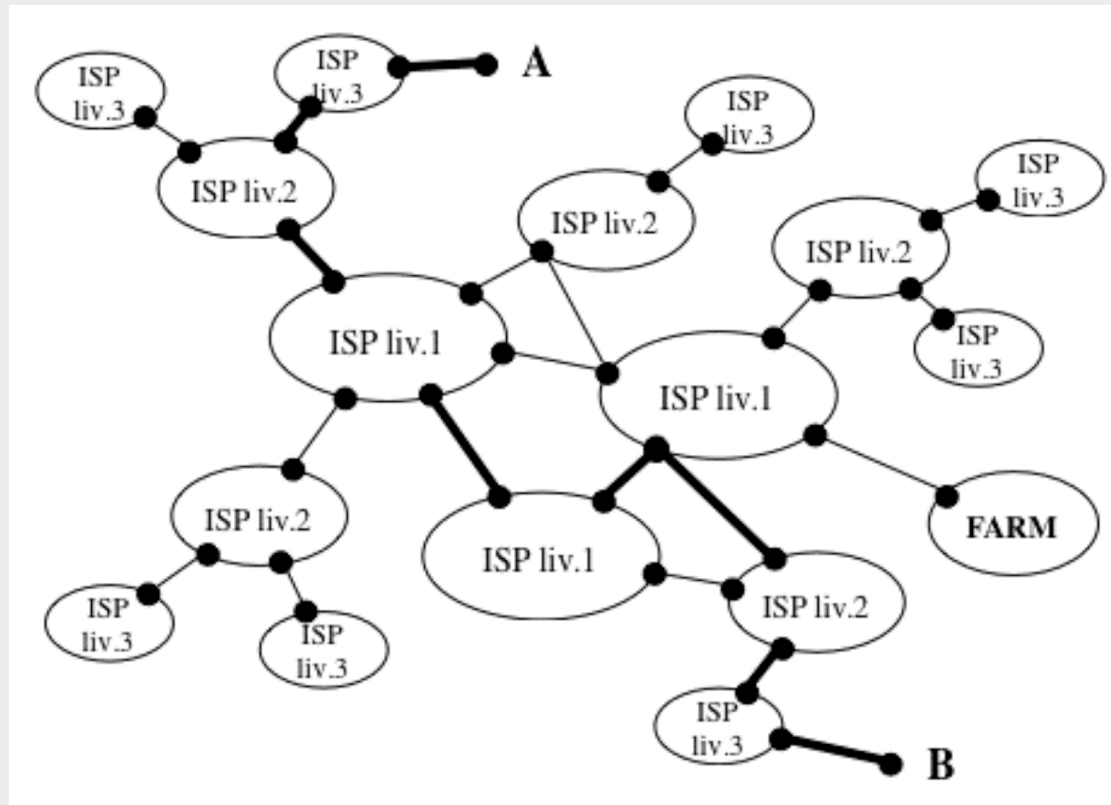
Internet as a graph

The central part of the net is composed of ISP's of various levels (tiers). Tier 1 includes the biggest ISP's, generally large communication companies or public infrastructures that compose the Internet backbone. Note that an ISP in tier 1 may have routers in different towns or countries, or even in different continents.

Medium size regional ISP's are found in tier 2 and directly connected with the backbone; still they may also be connected with some other ISP of the same tier.

Small local ISP's are similarly located in tier 3. Generally the hosts are connected to them, and the network traffic passes through local to regional providers, to the backbone, and back to regional and local providers, up to the final destination.

Internet as a graph



A, B, and FARM are clients. The routing path between A and B is indicated with an engrossed line. (From the lectures of R. Zheng).

Internet connection policy

For the connection, each client pays his provider, that in turn pays the next one, etc. In some cases the backbone providers (tier 1) do not charge directly for the connection but get the funding from their government. In many cases they have some control over the delivery of messages, e.g. blocking traffic from/to specific destinations. This is definitely true in China where all traffic is forced to pass through some control points, so a messages may take a very long path before reaching an albeit close destination.

In this organization, ISP's of tier 1 generally exchange messages without charging each other: in a sense, they operate as a unique huge backbone. This policy is peculiar of Internet, and may appear surprising because these ISP's are often big telecommunication companies that compete to attract the same clients. On the other hand, this behavior contributes to the overall efficiency of the system, then also to their own benefit.

Internet connection policy

Internet functioning is based on a mixture of cooperation and competition, two apparently contrasting words. To understand it, game theory and mathematical economics are often better tools than informatics. Let us just mention a study in this direction.

Games are simplified models of the real world. A well known example is the game of chicken (compared by Bertrand Russell to the proliferation of nuclear weapons):

	B: swerve	B: straight
A: swerve	0 , 0	-1 , +1
A: straight	+1 , -1	$-\infty$, $-\infty$

Nash Equilibrium

Nash Equilibrium

Internet connection policy

For example, routing protocols can be studied using game theory.

Packets are sent in groups. According to the most widely used strategy, if a group of a certain size has been successfully sent in a given time slot, the next group is increased by one packet; otherwise it is reduced to one half. It is believed that this corresponds to the equilibrium point of a (yet unknown) game.

In fact, the selfish strategy (sending each packet through the less crowded local link) has been mathematically analyzed as the equilibrium of a special game.

Internet: packet flow control

There are two levels of control, one for packet movement among AS's, and one for movement inside each AS. The former is guided by standard BGP (Border Gateway Protocol); the latter by an internal protocol of each AS, that may not be completely unveiled.

Although BGP should not include any control on the traffic flow, if not the actions for feeding the transmission, it allows a "manual reconfiguration" of the routers for allowing some AS decisions based on political, economic, or security reasons.

For example packets coming from a political organization may be denied transit, or be routed only to selected destinations.

Or the routers of a company may prevent traffic from reaching a competitor, or may allow transit only to paying clients.

In fact, nobody knows exactly what is going on.

Internet: the value of basic parameters

First note that Internet is a non directed net.

Second, keep in mind that Internet is continuously changing. Not only it has been growing constantly over the years, but large fluctuations occur at every minute, as the connections between AS's may be rearranged or may flicker. From time to time AS's with a few connections with the rest of the net may become completely isolated (invisible).


Internet constant growth means that the average number of newborn AS's is bigger than the one of AS's that go out of business. This latter number is surprisingly high (is in the order of one third of the number of newborn per year).

As Internet changes continuously, a discussion on its structure and the values of its parameters has only a meaning on the average.

Internet: basic parameters at AS level

Degree, betweenness, and shortest-path length distributions, together with the clustering coefficient, have been evaluated for Internet by several authors, both at AS level and at router level.

The raw data for these studies had been collected daily by the National Laboratory for Applied Network Research (NLNR) whose activity ended in June 2006.




Internet degree probability distribution at AS level, after Pastor-Satorras et al.

The power-law exponent is $\gamma = 2.2$

The average degree has been increasing over the years, and is now about 4.

Internet: basic parameters at AS level



Internet betweenness probability distribution at AS level, after Vázquez et al.


Recall that
$$\sigma(v) \equiv \sum_{i \neq j} \frac{B(i, v, j)}{B(i, j)}$$

N is the number of AS's

The power-law exponent γ is 2.1 ± 0.2

The clustering coefficient $C = \gamma / (z(z-1)/2)$ is about 0.2., hence much greater than the value of C for a random graph of comparable size.

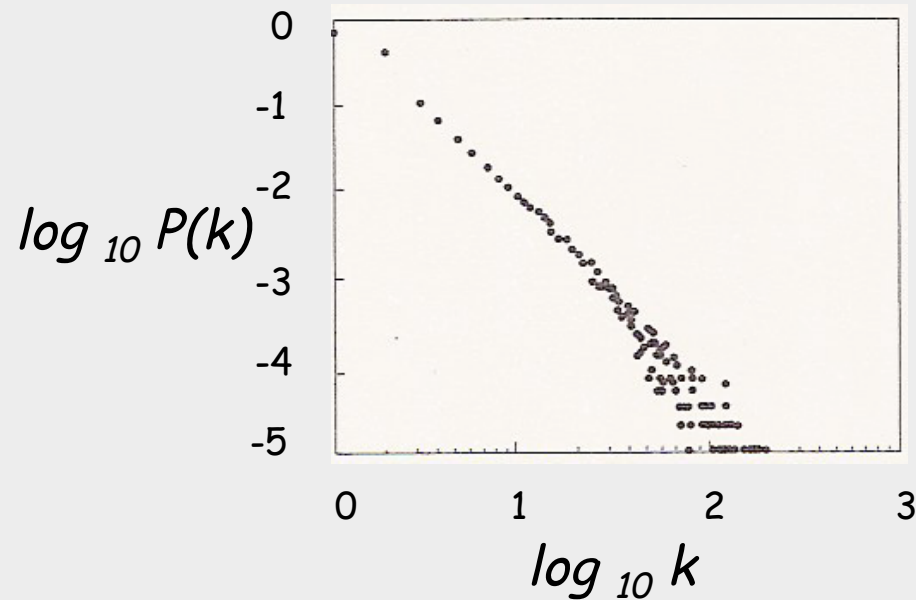
Internet: basic parameters at AS level



Internet shortest-path length distribution at AS level, after Vázquez et al.

The average shortest-path length is below 4, hence Internet IS A SMALL WORLD.

Internet: basic parameters at router level



Internet degree distribution at router level, after Govindan et al.

The power-law exponent γ is about 2.3

The average shortest-path length is about 10, not so far from the one of a random graph of comparable size. At a router level, Internet is greatly influenced by geographic factors (i.e., connections are much more frequent between close nodes).

WWW: birth and development

Born in Geneva in 1989, at CERN (Conseil Européen pour la Recherche Nucléaire).

The initial purpose was allowing all the researchers at CERN to browse through the lab project documents. Then it was made available outside, thus allowing easy document retrieval in Internet, once the "address" of a document is known.

Before the advent of WWW, searching through Internet was difficult and frustrating for the non expert. After WWW several browsers were born (Netscape, Explorer, Mozilla, etc.). Still, the real advancement for the majority of users occurred with the birth of search engines (Altavista, Yahoo!, Google, Ask, etc.).

WWW: basic concepts and terminology

WWW is an array of pages connected by links. The idea of hypertext (a text containing links to other texts) is crucial to the structure of WWW. Then WWW is a directed graph.

Although the pages of WWW reside in machines connected in Internet, there is in general no relation between the structure of Internet and the one of WWW. Two computers can be physically connected, but no WWW link may exist between them, while linked pages may reside in computer located very far. Moreover, copies of a WWW page may be present in many different machines.

The size (number of distinct pages) of WWW is not known. On the other hand it varies from minute to minute because a huge amount of pages are continuously built or deleted. The best estimate known, made by Gullì and Signorini, is >11.5 billion pages in January 2005.

Such an estimate can be made with mathematical techniques used by the people studying population growth.

WWW: estimating the size

Search engines declare the number of pages they are able to reach (in the order of billions). There is a consistent overlapping among those sets, while some pages are reachable only by one or some of the engines. To evaluate the (approximate) number of Web pages a sophisticated statistical technique is needed.

How to count the number T of orangutangs that live in a natural reserve in Borneo?


1. Capture M apes in day 1, mark them, and let them free.
2. Wait 30 days, during which you assume that the apes wander freely in the reserve. Then capture $N > M$ apes.
- 3 . Count the number Q of marked apes among the N . You may assume approximately that $T/M = N/Q$, i.e. $T = MN/Q$.

The structure of large growing graphs

A known phenomenon in graph theory is the appearance of a Giant Connected Component (*GCC*) in non equilibrium graphs built with random or preferential linking. *GCC* includes the vast majority of the N vertices as N grows, and appears if the edges are enough for the existence of a largest connected component with finite relative size R (i.e., R/N approaches a non-zero value). Then *GCC* grows with N .

If the graph is undirected, it consists of a growing *GCC* plus a set of finite independent connected components. If the graph is directed the situation is more complicated.

GCC in an directed graph



A directed graph with *GCC*, after Dorogovstev and Mendes (2003).


Consists of a *GWCC* (*W* for *weakly*, i.e., as if edges were undirected) containing a core *GSCC* (*S* for *strongly*, i.e., all vertices reachable through directed edges); Giant in- and out-component; and other secondary components.

According to a famous study of Broder et al., this is the structure of WWW.

Moreover the structure of *GSCC* is scalable, i.e., each (big enough) part of *GSCC* has the same structure of the whole graph. In fact, the distribution of such sub-graphs follow a power-law with exponent 2.5.

GCC in the Web

GSCC, *GIN-GSCC*, and *GOUT-GSCC* approximately contain 1/4 of the WWW vertices each.



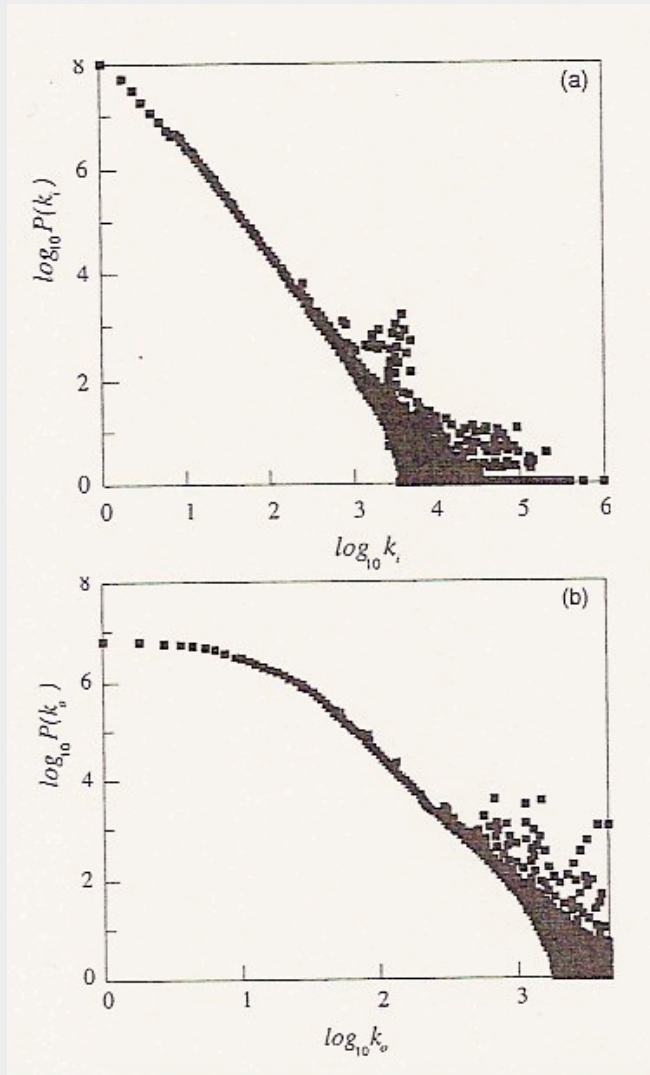
However, the structures of *GIN* and *GOUT* are very different:

From a vertex v of *GIN-GSCC* one can reach all the vertices in *GOUT*, but the average number of other vertices in *GIN* from which v can be reached is less than 200.

A vertex v of *GOUT-GSCC* can be reached by all the vertices in *GIN*, but the average number of other vertices in *GOUT* that can be reached from v is more than 3000.

This depends on how WWW is built (see later).

WWW: the value of basic parameters



Empirical in-degree (a) and out-degree (b) distributions of the whole Web, after Broder et al. (in Dorogovstev and Mendes, 2003).

In a first approximation, the power-law exponents are 2.1, and 2.7, respectively.

In a better approximations, $P(k)$ is proportional to $(k + c)^{-\gamma}$, with c and γ for in and out degree.

These values slightly differ for different families of Web pages. In particular, pages of scientific communities tend to have a much higher exponent than commercial ones.

WWW: the value of basic parameters

Several studies have shown experimentally that the WWW betweenness probability distribution follows a power-law.

Taking all edges as undirected, the average shortest path-length of WWW is about 7. So, the "undirected" WWW is a small world (a scarcely interesting result because Web links are directed).

In fact, the average directed shortest path-length of WWW is about 16. This does not mean that 16 clicks are required on average to pass from one page to another, because the user may not follow a shortest path.

There are pairs of pages separated by shortest paths of length > 1000 . Moreover, there are pairs of pages that cannot be reached from one another.

How does a new WWW page appear?

WWW basic parameters are well explained by the mechanisms of net growth/evolution. Consider three typical cases:

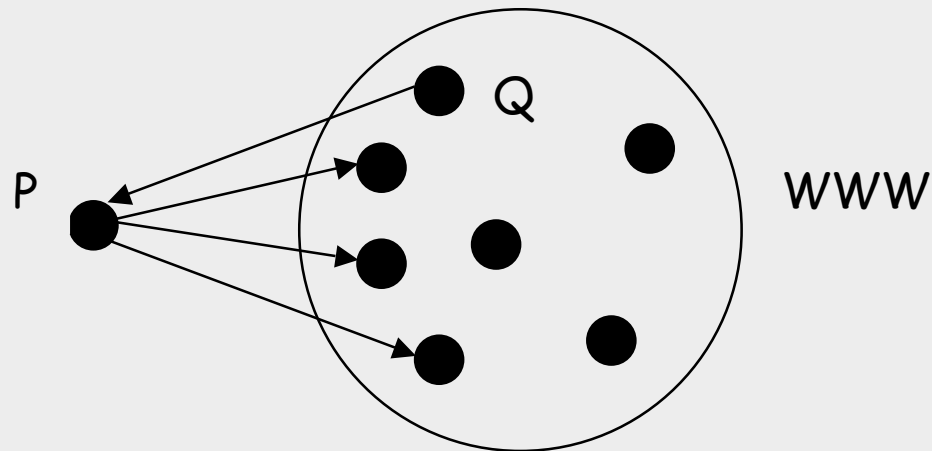
1. The simplest case. A new page P is created by somebody who has updating privileges on an already existing page Q (e.g., Q is a blog). P is first prepared (in general it contains links to other pages), then a link to P is inserted in Q . P is now on line.

2. A new page P is created within an institution or company C , e.g., the personal home page of a member of C . Page P is first prepared, then a link to P is inserted in the home page Q of C . P is now on line.

3. A page P already on line is modified, and its links to other pages are changed.

How does a new WWW page appear?

Cases 1 and 2 are as follows. Case 3 is less predictable.



The new edges from P to WWW obey preferential and probably random linking. Part of them may be directed to pages "geographically" close (e.g. in the same company site). Overall, a good mathematical model is PRDMM, yielding the power-law behavior observed in the experiments.

Internet names and addresses

Each Autonomous System (AS) in Internet is identified by a name, to which a numeric address corresponds.

For example the "Dipartimento di Informatica" of the "University of Pisa", in "Italy", has name di.unipi.it, and number 131.114.

Inside the AS further specifications occur. The router of the Dipartimento in charge of Web service has name www.di.unipi.it, and IP address 131.114.3.18. In fact this router is in subnetwork 3 of the Dipartimento, where is computer number 18.

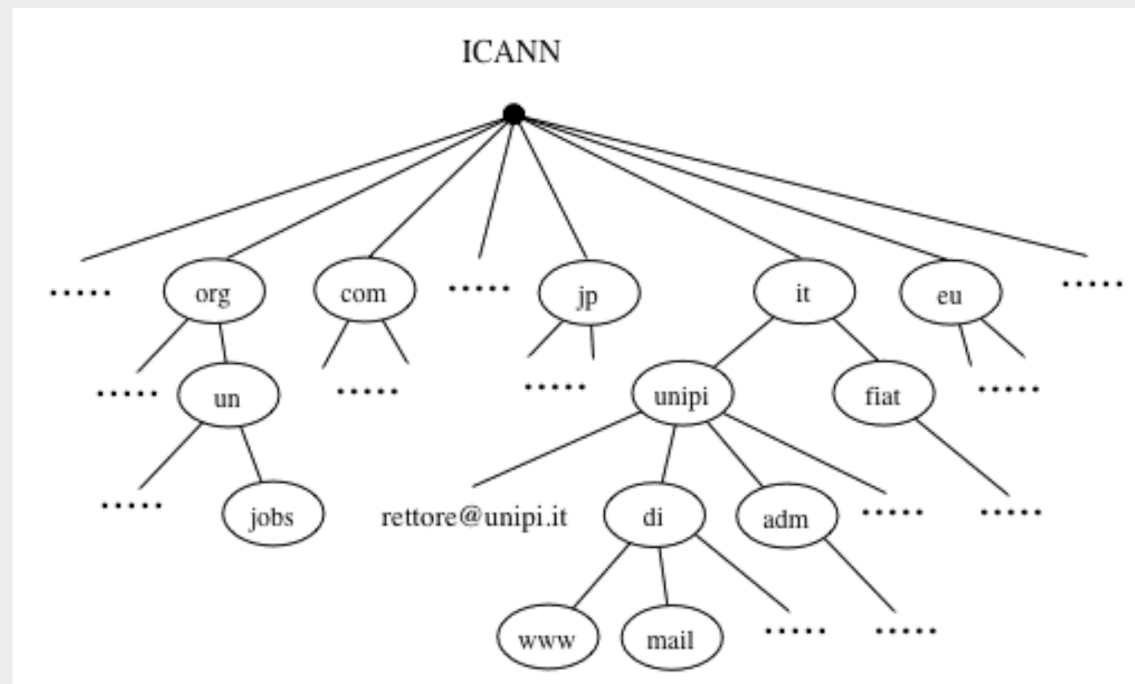
The structure of the IP addresses has some limited similarities with the telephone number system. The last digits of the IP address (3.18) denote a classification inside the AS. However, the first digits (131.114) do not denote a geographic or any other classification, as they are assigned according to availability (e.g., Italian names ending in .it do not necessarily contain 131 or 114). There is no structural dependency between names and addresses.

The tree of Domain Name System (DNS)

The set of Internet names has the structure of a tree.

The Dipartimento di Informatica (.di) is part of the University of Pisa (.unipi), that in turn is in Italy (.it).

A particular computer (mail) of the Dipartimento is in charge of e.mail service, for routing messages between Internet and the single computers. Another router (www) does Web service. All this is apparent in the Domain Names Tree:



DNS: birth and development

At the birth of Internet, all the information on site names and corresponding IP addresses was stored in a central computer at Stanford Research Institute (one of the first five ARPANET partners). Each night the users copied this ever changing archive.

With the growth of the archive, it was decided to structure it in tree form, assigning each tree node to a different institution in charge of maintaining consistency of the names of its children. The tables of correspondence between names and IP addresses are maintained in some nodes of the tree, so that retrieving the address corresponding to a given name entails traveling in the tree until a table with that information is found.

Each node has full control of its children. It may impose rules for joining the service (e.g., related to identification or fiscal regime), or deny access. On top, at the tree root, stands the Internet Corporation for Assigned Names and Numbers (ICANN).

ICANN: a political issue

ICANN is an American institution under the U.S. Department of Commerce. A big discussion is going on internationally on having ICANN in charge of Internet, as it could virtually impose any rule on its children, i.e. to the whole net or at least part of it.

Different countries, in particular Brazil, India, and, more mildly, the European Union, have been asking of passing the Internet control to ONU.

The American answer is many-sided.

First, the USA have sustained the birth and the major costs of Internet, so they claim the right of control.

Second, the majority of ICANN officers are not American, and have been operating in a democratic and transparent way.

Third, ONU has a ferocious bureaucracy, risking to paralyze the net.

The discussion is going on.

The problem of IP addresses

The present protocol IP_{v4} uses addresses of 32 bits, so in principle more than 4 billions addresses are possible. However, not all them are used, and availability of new addresses is becoming a serious problem.

Note that an IP address is a connection point to the net. The address may correspond to another net of arbitrary size, but this net has Internet access only through a specific router.

The protocol of the future IP_{v6} with 128 bits ($34 \cdot 10^{37}$ addresses) is completely defined, but is not yet in use because it would require big changes in the existing systems. Different countries, China in particular, are urging for the change.

DNS structure versus Internet

The names of DNS are linked through tree edges that show a strict relation between a name and its neighbors. However, there is no strict relation between the graph structure of Internet and the DNS tree.

A "domain" (e.g., .it) includes other sub-domains (.unipi.it), but there is no direct line between the servers of the two institutions that manage them, so that the physical connection may take place through a long chain of "hops".

Moreover, an institution in charge of a big domain (e.g., the Italian National Research Council is in charge of .it) may have many net routers placed in different locations or countries, to ease connection with its domain.

Summarizing on Internet, WWW, and DNS

Internet graph structure

- Net of AS each of which is a net, connected via peripheral routers. AS's are grouped in three tiers: tier 1 is the net backbone.
- Transmission based on packet switching via IP protocol.
- At AS level: power-law form of vertex degree and betweenness. Average shortest-path length about 4.
- At router level: power-law form of vertex degree and betweenness. Average shortest-path length about 10.
- Estimated number of participating computers: 1 billion in 2010.

Summarizing on Internet, WWW, and DNS

WWW graph structure

- Digraph with one *Giant Connected Component* about 1/4 of the total. This component is scalable.
- Power-law form of in- and out-degree, and betweenness. Average shortest-path length about 16.
- Crawling, indexing, and page ranking are crucial tasks of search engines.
- Estimated number of pages: 11.5 billions in 2005.

Summarizing on Internet, WWW, and DNS

DNS graph structure

- Tree of names corresponding to IP addresses.
- Each vertex (domain) is ruled by an institution put in charge of it by its parent vertex, with full control of children.
- ICANN at the tree root.

There is no relevant connection among the graph structures of Internet, WWW, and DNS.